

Deeply Coupled Cross-Modal Prompt Learning

Xuejing Liu¹, Wei Tang^{+ 2}, Jinghui Lu¹, Rui Zhao¹, Zhaojun Guo^{+ 3}, Fei Tan^{* 1}

¹ SenseTime Research

² Nanjing University of Science and Technology

³ Fudan University

{liuxuejing, lujinghui1, zhaorui, tanfei}@sensetime.com

weitang@njjust.edu.cn, zhaojun.guo1999@gmail.com

Abstract

Recent advancements in multimodal foundation models (e.g., CLIP) have excelled in zero-shot generalization. Prompt tuning involved in the knowledge transfer from foundation models to downstream tasks has gained significant attention recently. Existing prompt-tuning methods in cross-modal learning, however, either solely focus on language branch, or learn vision-language interaction in a shallow mechanism. In this context, we propose a **Deeply coupled Cross-modal Prompt learning (DCP)** method based on CLIP. DCP flexibly accommodates the interplay between vision and language with a **Cross-Modal Prompt Attention (CMPA)** mechanism, which enables the mutual exchange of respective representation through a well-connected multi-head attention module progressively and strongly. We then conduct comprehensive few-shot learning experiments on 11 image classification datasets and analyze the robustness to domain shift as well. Thorough experimental analysis evidently demonstrates the superb few-shot generalization and compelling domain adaption capacity of a well-executed DCP. The code can be found at <https://github.com/GingL/CMPA>.

1 Introduction

Large foundation models pre-trained on web-scale image-text pairs such as CLIP (Radford et al., 2021) and ALIGN (Jia et al., 2021) have shown promising performance on zero-shot image classification. Research has repeatedly shown that the general knowledge learned by the foundation models can also be transferred to diverse downstream tasks, such as few-shot image classification (Zhou et al., 2022b,a), visual grounding (Subramanian et al., 2022), visual question answering (Liu et al., 2022) and so on. They have exhibited a significant potential in open-vocabulary scenarios. Thus, the

challenge associated with how to efficiently and effectively adapt large pre-trained models to downstream tasks has garnered increasing attention especially in low-resource training scenarios.

Directly fine-tuning the foundation model is infeasible due to the massive training parameters and the catastrophic forgetting caused by overfitting (Kirkpatrick et al., 2016). In contrast, the parameter-efficient *prompt tuning* approach explored in natural language processing has yielded significant success (Lester et al., 2021), leading to an increased examination of this technique within the realm of multi-modality, especially in the language-branch of CLIP. For example, CoOp (Zhou et al., 2022b) and ProDA (Lu et al., 2022b) explore the vanilla few-shot learning based on CLIP by adjusting the embedding or distribution of the text prompt. CoCoOp (Zhou et al., 2022a) and ProGrad (Zhu et al., 2022) focus more on the unseen classes. They contextualize the text prompt either under the supervision of visual clues or tweak gradient direction to improve the generalization ability of the model.

The aforementioned approaches, however, only adjust the text embedding of CLIP and neglect the visual branch. The success of VPT (Jia et al., 2022) demonstrates the effectiveness of visual prompt learning. Inspired by this work, UPT (Zang et al., 2022) and MaPLe (Khattak et al., 2022) synergize the visual and textual prompts. Specifically, UPT improves the few-shot learning ability by generating visual and text prompts initially. MaPLe achieves better performance in the classification of unseen classes. They uncover the underlying rationale and limitations of dual-branch prompt tuning.

Concretely, the dual-branch CLIP learns the visual and language synergy only based on contrastive learning, whereas both branches lack mutual communication at the early stage of the network. Multi-modal prompt learning techniques,

⁺Work was done during internship at SenseTime Research

^{*}Corresponding author

such as MaPLe and UPT, incorporate language-vision interactions of the network and achieve substantially improved performance, highlighting the significance of the cross-modal interactions. However, previous studies have leveraged language-vision interactions at a superficial level. For example, UPT generates visual and text prompts before they are fed into the corresponding encoders. MaPLe generates visual prompts conditioned on language counterparts by a mapping function. Many studies (Dosovitskiy et al., 2021; Wang et al., 2022a) have shown that neural networks, especially transformer-based models, can leverage the deep fusion of information from multiple views to improve their performance. It remains less explored in the thread of multi-modal few-shot learning. To this end, we design **Deeply coupled Cross-modal Prompt learning (DCP)** enhancing the language-vision interaction. Specifically, DCP is built upon CLIP, with additional text and visual prompts across multiple layers. Different from previous methods with deep prompt tuning (Jia et al., 2022; Zang et al., 2022; Khattak et al., 2022), DCP only initializes the first layer of visual and text prompt randomly. The subsequent prompts are generated by **Cross-Modal Prompt Attention (CMPA)** module, which elegantly integrates the prompts from the preceding cross-modal layer. CMPA is characterized with stronger connection in two folds, i.e., *Depth* and *Breadth*. 1) *Depth* means that CMPA intensifies the correlation of the prompts among different layers. 2) *Breadth* refers to that CMPA amplifies the interaction between visual and language modalities. CMPA is the core module to realize the deep coupling between two modalities. Essentially, DCP empowered by CMPA amalgamates uni-branch and dual-branch multi-modal pre-training paradigms in a favorable way in an attempt to bridge the discrepancy between visual and textual knowledge without introducing too much overhead.

To conclude, the contributions of this work are as follows:

- We develop a deeply coupled cross-modal prompt learning (DCP) method with a core module cross-modal prompt attention (CMPA). CMPA can reinforce the interaction between visual and language modals across different layers.
- We benchmark our method on 11 image classification datasets consisting of generic objects,

scenes, actions and fine-grained categories. Our method surpasses visual prompt tuning, text prompt tuning and existing competitive multi-modal prompt tuning methods under the few-shot setting.

- We conduct experiments on domain adaptation tasks. Our method achieves comparable performance to the state-of-the-art methods, indicating the robustness of our method to domain shift.

2 Related Work

2.1 Vision-language Pre-trained Models

The advent of Transformer (Vaswani et al., 2017) has accelerated the development of large-scale pre-training. The application of Transformer in the multi-modal is divided into two schools of thought: one is the single-stream model, in which language and vision information are fused at the beginning and fed directly into the encoder together; the other is the dual-stream model, in which language and vision information first pass through two separate encoder modules at the beginning, and then the different modal information is fused through the cross Transformer.

At the outset, the basic architecture of some contemporaneous work is BERT. The images are detected with Faster-RCNN (Ren et al., 2015) for region features, and these image region features are fed into BERT along with text information to align the text and image information. Following the same process as BERT, these methods first pre-train and then fine-tune on the corresponding tasks. Single-stream networks (Li et al., 2019; Alberti et al., 2019; Chen et al., 2019; Li et al., 2020; Su et al., 2020; Zhou et al., 2020; Qi et al., 2020; Lu et al., 2020) fuse information from different modalities directly through an encoder. The dual-stream models (Lu et al., 2019; Tan and Bansal, 2019) integrate different modal information through cross modal transformer. Empirically single-stream networks are more sufficient for information fusion, while dual-stream networks can be more efficient for training due to fewer training parameters. In the design of our method, we aim to combine the advantages of the single-stream and dual-stream, so as to enhance the cross-modal integration without introducing many training parameters.

Recent cross-modal large-scale pre-training models have made greater breakthroughs in train-

ing data scale and tasks by devising various model architectures and training objectives, and have achieved impressive performance in many downstream tasks. CLIP (Radford et al., 2021) and ALIGN (Jia et al., 2021) got remarkable zero-shot results after being pre-trained on millions or billions of (image, text) pairs collected from the internet. Coca (Yu et al., 2022) combined the advantages of the contrast learning method (Radford et al., 2021) and the generative model SiMVLN (Wang et al., 2022b) by adding caption loss to the contrast loss of CLIP. OFA (Wang et al., 2022a), Unified-IO (Lu et al., 2022a) and Florence (Yuan et al., 2021) unified vision, language and multi-modal tasks by pre-training on both cross-modal and uni-modal data. These methods have achieved state-of-the-art results in many downstream tasks. Some methods are dedicated to improving the performance of certain specific tasks. UniTAB (Yang et al., 2022) focused on grounded vision-language tasks such as grounded captioning and visual grounding. GLIP (Li et al., 2022) unified object detection and phrase grounding for pre-training. Pre-training models have opened up a situation where deep learning models scale and perform in tandem, becoming a revolutionary breakthrough in artificial intelligence and deep learning.

2.2 Prompt Learning

For a long time, first pre-training then fine-tuning was the dominant approach to apply large foundation models to downstream tasks. However, fine-tuning for large models is inefficient and may cause catastrophic forgetting (Kirkpatrick et al., 2016). Prompt learning is proposed to address the above problems. The prompt is usually a series of trainable parameters inserted into the input. The success of prompt learning in NLP (Lester et al., 2021) has inspired its application in other modalities. VPT (Jia et al., 2022) is a typical successful application of prompt learning on computer vision. Prompt learning has generated more attention and made great progress in cross-modal learning.

SoftCPT (Ding et al., 2022) and CPL (He et al., 2022) applied prompt tuning to different vision and language tasks and outperformed single-task prompt tuning method. CoOp (Zhou et al., 2022b), ProDA (Lu et al., 2022b) and UPT (Zang et al., 2022) adapted prompt learning to traditional few-shot visual recognition with CLIP as the backbone.

CoCoOp (Zhou et al., 2022a), ProGrad (Zhu et al., 2022) and MaPLe (Khattak et al., 2022) improved the classification performance of pre-trained models on novel categories by prompt learning. Different from previous methods, our approach brings stronger connection between modalities and layers with proposed cross-modal prompt attention. The stronger interaction between vision and language enables our method to get state-of-the-art performance in the few-shot learning.

3 Method

In this section, we first introduce the preliminaries, including CLIP (Radford et al., 2021), CoOp (Zhou et al., 2022b) and VPT (Jia et al., 2022). Then, we describe our deeply coupled prompt learning (DCP) and detail its underlying module CMPA.

3.1 Preliminaries

CLIP is a dual-encoder pre-trained model which consists of a text encoder and an image encoder. The text and image are independently encoded by the corresponding encoder, then projected to the same embedding space by a projection layer. Specifically, the backbone of the image encoder is ResNet (He et al., 2016) ($d=256$) or ViT ($d=512$), which can map the high-dimension image into a low-dimension embedding. The text encoder is built based on the decoder of Transformer (Vaswani et al., 2017), which is also known as GPT (Brown et al., 2020), to generate a vectorized representation for a sequence of words. The model uses a contrastive loss to align the two modalities during training stage. The training objective is to maximize the cosine similarity for the match image-text pairs and minimize the unmatched ones.

In zero-shot image recognition, the image encoder of CLIP encodes the image into a feature representation \mathbf{x} . The input text is usually in the form of “a photo of a {class}.” (*discrete prompt*), where the “{class}” token is the name of each category. For each dataset containing K categories, a set of text prompts $\{\mathbf{w}_i\}_{i=1}^K$ are generated by the text encoder. The prediction probability is computed as

$$p(y | \mathbf{x}) = \frac{\exp(\cos(\mathbf{x}, \mathbf{w}_y) / \tau)}{\sum_{i=1}^K \exp(\cos(\mathbf{x}, \mathbf{w}_i) / \tau)}, \quad (1)$$

where τ is a temperature parameter.

CoOp adapts CLIP to downstream tasks with prompt tuning. Specifically, CoOp tries to learn

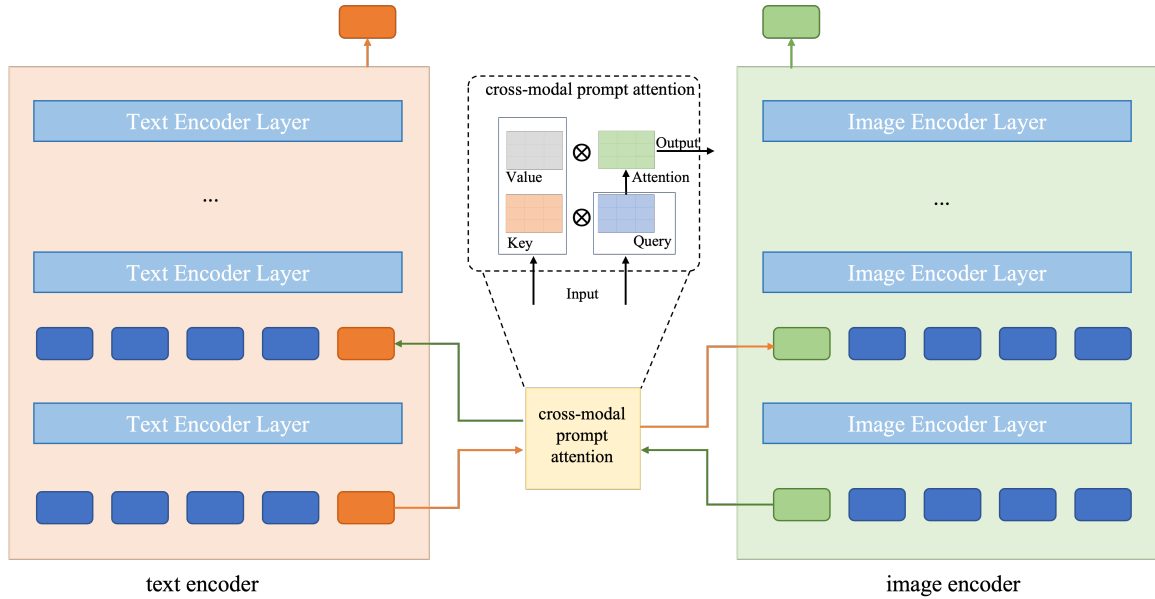


Figure 1: The architecture of deeply coupled prompt learning and cross-modal prompt attention module.

prompt embedding (*continuous prompt*) during few-shot training to avoid manual prompts. The prompt fed in the text encoder is designed as $t = [V]_1[V]_2\dots[V]_M[CLASS]$, where $[V]_m$ ($m \in \{1, \dots, M\}$) is initialized with the same dimension as word embeddings. The parameters of the CLIP model is frozen while the prompt is trainable. The prediction probability of CoOp is

$$p(y | \mathbf{x}) = \frac{\exp(\cos(\mathbf{x}, g(\mathbf{t}_y))/\tau)}{\sum_{i=1}^K \exp(\cos(\mathbf{x}, g(\mathbf{t}_i))/\tau)}, \quad (2)$$

where $g(\cdot)$ denotes the text encoder.

VPT is an efficient and effective way to adapt large-scale Transformer models in vision with only a small amount of trainable parameters. The backbone of VPT is ViT, which is the same as the image encoder of CLIP. There are two variants of VPT: VPT-Shallow and VPT-Deep. VPT-Shallow only inserts prompts into the first layer of the Transformer. The visual prompt can be defined as $p = [P]_1[P]_2\dots[P]_N$, where $[P]_n$ ($n \in \{1, \dots, N\}$) keeps the same dimension as the image embedding. The input of VPT-shallow is $[x_{cls}, p, x]$, where x_{cls} is the classification token $[CLS]$. VPT-Deep introduces visual prompts at every Transformer layer. The deep VPT can be formulated as

$$\begin{aligned} [\mathbf{x}_{cls}^i, \dots, \mathbf{x}^i] &= L^i([\mathbf{x}_{cls}^{i-1}, \mathbf{p}^{i-1}, \mathbf{x}^{i-1}]) \\ i &= 1, 2, \dots, L \\ \mathbf{y} &= \text{Head}(\mathbf{x}_{cls}^L), \end{aligned} \quad (3)$$

where L denotes the number of Transformer layers and $Head$ is the classification head. Only the prompts and classification head is learnt during training. VPT achieves impressive performance on 24 downstream recognition tasks.

3.2 Cross-modal Prompt Attention

Inspired by the advance of prompt learning in vision and language, recent studies start to explore multi-modal prompt learning (Zang et al., 2022; Khattak et al., 2022). These methods update the visual and text prompt simultaneously to achieve balance in the learning of visual and text embedding. Although the visual and text embedding are adapted to the few-shot data, the interaction between visual and text is still insufficient. Hence we propose deeply coupled cross-modal prompt learning (DCP), which can enhance the communication between prompts across different layers and modalities. The essential module of DCP is cross-modal prompt attention, which fuses visual and text with multi-head cross-modal attention. Figure 1 depicts the pipeline of DCP and the detailed architecture of cross-modal prompt attention (CMPA).

Our method follows the implementation of CLIP, which is also a dual-encoder model. Differently, we add prompts to every branch, and enable information fusion between vision and language during training through CMPA. Specifically, CMPA is a multi-head attention with visual and text prompts as inputs. The language prompts of the first layer are initialized with the pre-trained CLIP word em-

beddings of the template 'a photo of a <class>', whereas the visual prompts inserted into the first layer are randomly initialized from a normal distribution. Then, the prompts of the next layer are generated by CMPA based on the prompts from the preceding layer. Formally, CMPA can be formulated as

$$P_t^{l+1} = \text{softmax} \left(\frac{P_v^l (P_t^l)^T}{\sqrt{d_k}} \right) P_t^l \quad (4)$$

$$P_v^{l+1} = \text{softmax} \left(\frac{P_t^l (P_v^l)^T}{\sqrt{d_k}} \right) P_v^l \quad (5)$$

$$l = 1, 2, \dots, N - 1, \quad (6)$$

where P_t^l and P_v^l denote the text prompt and visual prompt the l layer of each encoder, respectively. N is the depth of CMPA, which is smaller than the length of text and visual encoder. d_k is the dimension of keys.

Different from previous methods, only the prompts from the first layer are randomly generated. The subsequent prompts condition on the prompts from both visual and language modal. CMPA enables information communication between vision and text through corresponding prompts. Totally, CMPA brings stronger feature fusion from two aspects: layers and modalities. Note that CMPA shares parameters from different layers, and the additional trainable parameters is only in a small amount.

4 Experiments

In this section, we conduct experiments to evaluate the effectiveness of our method under two settings. One is few-shot visual recognition including 11 different datasets covering generic objects, scenes, actions and fine-grained categories. The other is domain adaptation, where we train our model on ImageNet and evaluate it on other four datasets.

4.1 Few-shot Learning

4.1.1 Datasets

Following CoOp (Lester et al., 2021), we evaluate our method on 11 public visual recognition datasets: ImageNet (Deng et al., 2009), Caltech101 (Fei-Fei et al., 2004), OxfordPets (Parkhi et al., 2012), StanfordCars (Krause et al., 2013), Flowers102 (Nilsback and Zisserman, 2008), Food101 (Bossard et al., 2014), FGVCaircraft (Maji et al., 2013), SUN397 (Xiao et al.,

2010), DTD (Cimpoi et al., 2014), EuroSAT (Helber et al., 2019) and UCF101 (Soomro et al., 2012). We also use the same 1, 2, 4, 8 and 16 shots as CoOp for training and the full test set for evaluation purpose. The reported results are the average over three runs with different random seeds.

4.1.2 Implementation Details

We use the pre-trained ViT-B/16 CLIP model as our backbone. The length of prompt tokens for visual and textual context are both 16. The prompt depth is 9 as a trade-off between accuracy and training efficiency. We set the batch-size to 4 with a learning rate of 0.0035 via SGD optimizer. We use 20 epochs for most datasets, except ImageNet, SUN397 and Food101. Also, 5-epoch setting works for diverse shots of Food101, 1/2/4-shot of ImageNet, and 1/2-shot of SUN397, respectively.

4.1.3 Main Results

Baseline Methods. We compare our method with the original zero-shot CLIP, text prompt learning (CoOp), visual prompt learning (VPT) and multi-modal prompt learning (MaPLe), which all have ViT-B/16 as visual backbone. Basically, we follow the implementation of MaPLe (Khattak et al., 2022). The prompt length of CoOp is set to 16. VPT uses a prompt length of 8 and the visual and text prompt length of MaPLe is 2. The training epoch of CoOp is defined as 10, and that of VPT and MaPLe is 5. We use the deep variant of VPT in few-shot experiments. The prompt depth of MaPLe is 9 as their original setting.

Performance Analysis. Figure 2 demonstrates our results comparison with other methods. The top left sub-figure shows the average performance of four methods. We can have the following findings. 1) Overall, cross-modal prompt learning (DCP and MaPLe) gets a large performance gain compared with single-modal prompt learning methods (VPT and CoOp). VPT and CoOp achieve comparable performance on different shots. These results demonstrate the superiority of cross-modal prompt learning over uni-modal prompt learning. 2) Although both belong to multi-modal prompt learning methods, our method still outperforms MaPLe on 1/2/4/8/16 shots settings by 1.72/3.18/3.19/2.20/2.76(%). MaPLe utilized a linear layer to generate visual prompts from text prompts. Our proposed DCP enhances the interaction between vision and language with a cross-

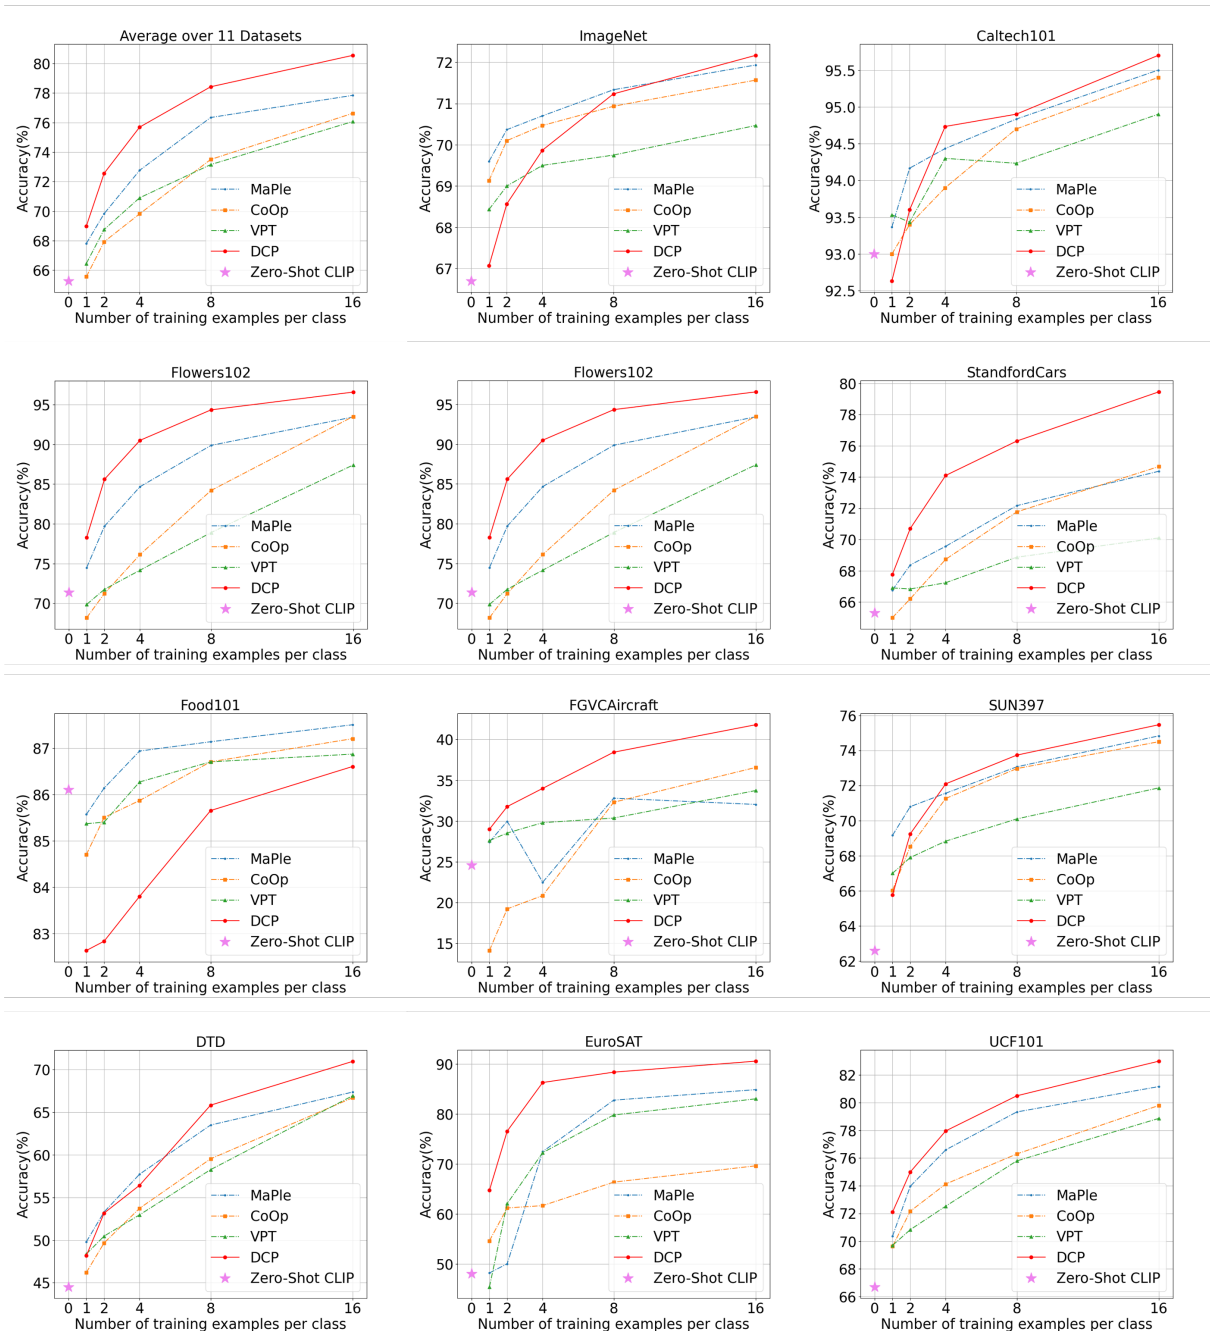


Figure 2: **Main results of few-shot image classification on 11 datasets.** The accuracy (%) is the average over three runs on 1/2/4/8/16 shots. Overall, our DCP (red line) outperforms other methods by a large margin on the average results of 11 datasets.

modal prompt attention, which can not only guide visual embedding learning through text prompts, but also influence the language embedding with visual prompts. 3) Compared with 2/4/8/16 shots, our approach achieves a lower performance gain on one shot. We can also find that on separate datasets, our method achieves the best performance in almost all 16-shot cases (except for Food101). This phenomenon indicates that our method is more effective in cases where the number of shots is rela-

tively large. This is probably because the alignment between different modals is more challenging due to the small number of samples per category.

For individual datasets, we find that our approach has significant performance improvements on Flowers102, StanfordCars, FGVCaircraft, and EuroSAT. However, on the datasets of general categories such as ImageNet and Caltech101, our method does not achieve satisfactory performance when the number of shots is less than 16. We can

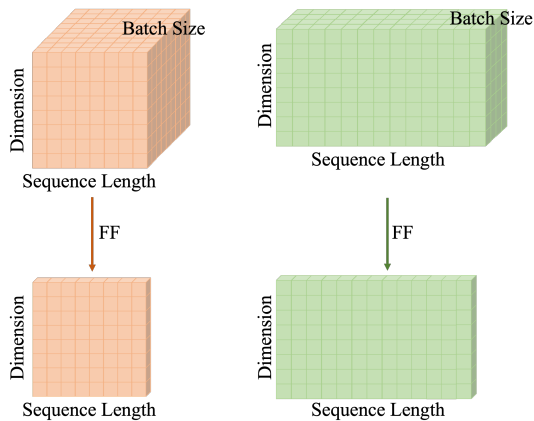


Figure 3: The illustration of feature fusion (FF). The left branch represents the text prompt, and the right shows the visual prompt.

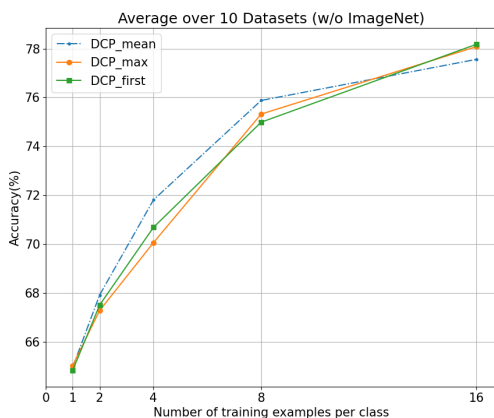


Figure 4: The comparison of different feature fusion methods on 10 datasets without ImageNet.

conclude that our method is more robust for fine-grained classification datasets, and we need more shots for general category classification. On the dataset of Food101, our method performs slightly lower than MaPLe. We also find that all methods underperform zero-shot on 1-shot setting. We suppose this phenomenon comes from the noisy training data of Food101 (Bossard et al., 2014).

4.1.4 Ablation Study

There are two important settings in CMPA: the feature fusion method in different prompts and parameter sharing of CMPA across different layers. We conduct corresponding ablation experiments in this section to find the optimal setting.

Feature Fusion in Prompts. Before the visual and text prompts are fed into the CMPA, the dimension of the batch size is supposed to be consistent.

Variant	2	4	6	8	16
w/ PS	68.99	72.56	75.69	78.42	80.55
w/o PS	67.42	71.34	75.27	78.49	80.53

Table 1: The performance comparison with and without parameter sharing. The results are the average accuracy on 11 datasets of different shots.

The defined batch size only affects visual prompt while the batch size of text prompts is actually the number of the dataset due to the implementation of CLIP. The dimension transformation of visual and text prompts is shown in Figure 3. The batch size of text prompt is actually the number of categories in the dataset. We experiment with three settings to align the batch size of visual and text prompts. Figure 4 reports the average accuracy over three runs on different shots (1/2/4/8/16) of 10 datasets (without ImageNet for time efficiency). ‘Avg’ means that we use the average of visual and text prompts across the dimension of batch. ‘Max’ stands for using the features with the highest response across the batch dimension as the visual and text prompt. ‘First’ represents that we select the first embedding across the batch dimension of visual and text prompts to feed into CMPA. Overall, the ‘avg’ setting of feature fusion can achieve better performance compared with ‘max’ and ‘first’.

Parameter Sharing. We intend to learn as few parameters as possible to achieve a transfer of large-scale pre-trained models in downstream tasks. Setting the prompt depth to 9 means that there are 9 CMPA modules, which greatly increases the number of trainable parameters for the model. Hence we conduct the experiment in which the parameters of CMPA are shared across different layers. Table 1 shows the average results of different shots on 11 datasets. ‘PS’ is short for ‘parameter sharing’. It can be observed that on most shots (except for 8 shots) the performance of parameter sharing is higher than non-sharing setting.

4.2 Domain Generalization

After prompt tuning on specific datasets, we do not want to lose the general knowledge of the pre-trained large model. In this section, we conduct domain adaptation experiments to evaluate the generalization ability of our model DCP.

Method	Source		Target			Average	OOD Average
	ImageNet	-V2	-S	-A	-R		
CLIP	66.73	60.83	46.15	47.77	73.96	59.09	57.18
CoOp	71.53	64.20	47.99	49.71	75.21	61.73	59.28
CoCoOp	71.02	64.07	48.75	50.63	76.18	62.13	59.91
VPT-Deep	70.57	63.67	47.66	43.85	74.42	60.03	57.40
MaPLe	71.02	64.07	49.15	50.90	76.98	<u>62.42</u>	60.28
UPT	72.63	<u>64.35</u>	48.66	<u>50.66</u>	76.24	62.51	<u>59.98</u>
DCP (ours)	<u>71.53</u>	64.50	<u>48.77</u>	49.40	<u>76.50</u>	62.14	59.79

Table 2: Domain generalization comparison of DCP with existing approaches. The winners and runners-up are marked in bold font and underlined, respectively.

4.2.1 Datasets and Implementation Details

Following (Zhou et al., 2022b), we use ImageNet (Deng et al., 2009) as source domain, and ImageNet V2 (Recht et al., 2019), ImageNet-Sketch (Wang et al., 2019), ImageNet-A (Hendrycks et al., 2021b) and ImageNet-R (Hendrycks et al., 2021a) as target domains. We train our model on the 16 shots of ImageNet, and test it on other four datasets. Different from the settings in few-shot task, the training epoch on 16-shot ImageNet in cross domain task is set to 5. We also decrease the prompt length to 8.

4.2.2 Main Results

Table 2 compares our method DCP with other prompt learning methods on cross-domain tasks. The compared methods include zero-shot CLIP, unimodal prompt learning methods (CoOp, CoCoOp and VPT-Deep) and multi-modal prompt learning methods (MaPLe and UPT). The best results on different datasets are in bold, and the second best results are underlined. We can observe that 1) prompt learning does not corrupt the generalization ability of pre-trained large models; 2) multi-modal prompt learning methods outperform unimodal prompt learning methods in generalization performance; 3) our method can get comparable performance as the state-of-the-art methods.

5 Discussion and Conclusion

This paper proposes a deeply coupled cross-modal prompt learning method, with a core module cross-modal prompt attention. Our method focuses on optimizing the interaction across different models and layers to address the alignment between vision and language. Experiments on few-shot image clas-

sification and domain adaptation evidence that our method can transfer the general knowledge learned by pre-trained foundation models to downstream tasks without penalty of the original generalization ability. Our method provides a strong baseline on few-shot image classification. The deep fusion between visual and language information may enable our approach to have greater potential for complex cross-modal tasks, such as referring expression comprehension (Subramanian et al., 2022), image retrieval (Baldrati et al., 2022) and visual question answering (Liu et al., 2022). We will apply our method to such complicated cross-modal tasks to evaluate its effectiveness in our future work.




Figure 5: Accuracy comparison of different epochs on Sun397 and ImageNet.

6 Limitations

We discover that for datasets with a relatively large number of categories, our method requires a more delicate setting of epoch under different shots. Figure 5 shows the average results on Sun397 and ImageNet of different epochs. It can be observed that for datasets with a large number of categories (such as Sun397 and ImageNet), as the number of shots decreases, the performance deteriorates with an increase in the number of epochs, which is not evident on the datasets with a small number of cat-

egories. We will delve further into this problem to find the reason and solution.

7 Acknowledgement

We would like to thank anonymous reviewers for their insightful comments to help improve the paper. This publication has emanated from research conducted with the support of SenseTime Research and Hetao Shenzhen-Hong Kong Science and Technology Innovation Cooperation Zone (HZQB-KCZYZ-2021045).

References

- Chris Alberti, Jeffrey Ling, Michael Collins, and David Reitter. 2019. [Fusion of detected objects in text for visual question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 2131–2140. Association for Computational Linguistics.
- Alberto Baldrati, Marco Bertini, Tiberio Uricchio, and Alberto Del Bimbo. 2022. [Effective conditioned and composed image retrieval combining clip-based features](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 21434–21442. IEEE.
- Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. 2014. Food-101 - mining discriminative components with random forests. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VI*, volume 8694 of *Lecture Notes in Computer Science*, pages 446–461. Springer.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. [UNITER: learning universal image-text representations](#). *CoRR*, abs/1909.11740.
- Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. 2014. [Describing textures in the wild](#). In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 3606–3613. IEEE Computer Society.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. [Imagenet: A large-scale hierarchical image database](#). In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pages 248–255. IEEE Computer Society.
- Kun Ding, Ying Wang, Pengzhang Liu, Qiang Yu, Haojian Zhang, Shiming Xiang, and Chunhong Pan. 2022. [Prompt tuning with soft context sharing for vision-language models](#). *CoRR*, abs/2208.13474.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Li Fei-Fei, Rob Fergus, and Pietro Perona. 2004. [Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories](#). In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2004, Washington, DC, USA, June 27 - July 2, 2004*, page 178. IEEE Computer Society.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- Xuehai He, Diji Yang, Weixi Feng, Tsu-Jui Fu, Arjun R. Akula, Varun Jampani, Pradyumna Narayana, Sugato Basu, William Yang Wang, and Xin Eric Wang. 2022. [CPL: counterfactual prompt learning for vision and language models](#). *CoRR*, abs/2210.10362.
- Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. 2019. [Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification](#). *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.*, 12(7):2217–2226.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. 2021a. [The many faces of robustness: A critical analysis of out-of-distribution generalization](#). In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 8320–8329. IEEE.

- Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. 2021b. [Natural adversarial examples](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 15262–15271. Computer Vision Foundation / IEEE.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. [Scaling up visual and vision-language representation learning with noisy text supervision](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 4904–4916. PMLR.
- Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge J. Belongie, Bharath Hariharan, and Ser-Nam Lim. 2022. [Visual prompt tuning](#). In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXIII*, volume 13693 of *Lecture Notes in Computer Science*, pages 709–727. Springer.
- Muhammad Uzair Khattak, Hanoona Abdul Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. 2022. [Maple: Multi-modal prompt learning](#). *CoRR*, abs/2210.03117.
- James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2016. [Overcoming catastrophic forgetting in neural networks](#). *CoRR*, abs/1612.00796.
- Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 2013. [3d object representations for fine-grained categorization](#). In *2013 IEEE International Conference on Computer Vision Workshops, ICCV Workshops 2013, Sydney, Australia, December 1-8, 2013*, pages 554–561. IEEE Computer Society.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 3045–3059. Association for Computational Linguistics.
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020. [Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 11336–11344. AAAI Press.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. [Visualbert: A simple and performant baseline for vision and language](#). *CoRR*, abs/1908.03557.
- Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and Jianfeng Gao. 2022. [Grounded language-image pre-training](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 10955–10965. IEEE.
- Yuhang Liu, Wei Wei, Daowan Peng, and Feida Zhu. 2022. [Declaration-based prompt tuning for visual question answering](#). In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 3264–3270. ijcai.org.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. [Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13–23.
- Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Mottaghi, and Aniruddha Kembhavi. 2022a. [Unified-io: A unified model for vision, language, and multi-modal tasks](#). *CoRR*, abs/2206.08916.
- Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 2020. [12-in-1: Multi-task vision and language representation learning](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10434–10443. Computer Vision Foundation / IEEE.
- Yuning Lu, Jianzhuang Liu, Yonggang Zhang, Yajing Liu, and Xinmei Tian. 2022b. [Prompt distribution learning](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 5196–5205. IEEE.
- Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew B. Blaschko, and Andrea Vedaldi. 2013. [Fine-grained visual classification of aircraft](#). *CoRR*, abs/1306.5151.
- Maria-Elena Nilsback and Andrew Zisserman. 2008. [Automated flower classification over a large number of classes](#). In *Sixth Indian Conference on Computer Vision, Graphics & Image Processing, ICVGIP 2008, Bhubaneswar, India, 16-19 December 2008*, pages 722–729. IEEE Computer Society.
- Omkar M. Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. 2012. [Cats and dogs](#). In *2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012*, pages 3498–3505. IEEE Computer Society.

- Di Qi, Lin Su, Jia Song, Edward Cui, Taroon Bharti, and Arun Sacheti. 2020. [Imagebert: Cross-modal pre-training with large-scale weak-supervised image-text data](#). *CoRR*, abs/2001.07966.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. 2019. [Do imagenet classifiers generalize to imagenet?](#) In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 5389–5400. PMLR.
- Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. [Faster R-CNN: towards real-time object detection with region proposal networks](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 91–99.
- Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. 2012. [UCF101: A dataset of 101 human actions classes from videos in the wild](#). *CoRR*, abs/1212.0402.
- Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2020. [VL-BERT: pre-training of generic visual-linguistic representations](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Sanjay Subramanian, William Merrill, Trevor Darrell, Matt Gardner, Sameer Singh, and Anna Rohrbach. 2022. [Reclip: A strong zero-shot baseline for referring expression comprehension](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 5198–5215. Association for Computational Linguistics.
- Hao Tan and Mohit Bansal. 2019. [LXMERT: learning cross-modality encoder representations from transformers](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 5099–5110. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Haohan Wang, Songwei Ge, Zachary C. Lipton, and Eric P. Xing. 2019. [Learning robust global representations by penalizing local predictive power](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 10506–10518.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022a. [OFA: unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 23318–23340. PMLR.
- Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. 2022b. [Simvlm: Simple visual language model pretraining with weak supervision](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Jianxiong Xiao, James Hays, Krista A. Ehinger, Aude Oliva, and Antonio Torralba. 2010. [SUN database: Large-scale scene recognition from abbey to zoo](#). In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2010, San Francisco, CA, USA, 13-18 June 2010*, pages 3485–3492. IEEE Computer Society.
- Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Faisal Ahmed, Zicheng Liu, Yumao Lu, and Lijuan Wang. 2022. [Unitab: Unifying text and box outputs for grounded vision-language modeling](#). In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXVI*, volume 13696 of *Lecture Notes in Computer Science*, pages 521–539. Springer.
- Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. 2022. [Coca: Contrastive captioners are image-text foundation models](#). *CoRR*, abs/2205.01917.
- Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, Ce Liu, Mengchen Liu, Zicheng Liu, Yumao Lu, Yu Shi, Lijuan Wang, Jianfeng Wang, Bin Xiao, Zhen Xiao, Jianwei Yang, Michael Zeng, Luwei Zhou, and Pengchuan Zhang. 2021. [Florence: A new foundation model for computer vision](#). *CoRR*, abs/2111.11432.
- Yuhang Zang, Wei Li, Kaiyang Zhou, Chen Huang, and Chen Change Loy. 2022. [Unified vision and language prompt learning](#). *CoRR*, abs/2210.07225.

- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022a. [Conditional prompt learning for vision-language models](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 16795–16804. IEEE.
- Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. 2022b. [Learning to prompt for vision-language models](#). *Int. J. Comput. Vis.*, 130(9):2337–2348.
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J. Corso, and Jianfeng Gao. 2020. [Unified vision-language pre-training for image captioning and VQA](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 13041–13049. AAAI Press.
- Beier Zhu, Yulei Niu, Yucheng Han, Yue Wu, and Hanwang Zhang. 2022. [Prompt-aligned gradient for prompt tuning](#). *CoRR*, abs/2205.14865.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 6
- A2. Did you discuss any potential risks of your work?
Section 6
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
No response.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
No response.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
No response.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
No response.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
No response.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
No response.

C Did you run computational experiments?

Section 4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Left blank.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

section 4.1.2

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

section 4.1.1

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Left blank.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.